



Kenneth L. Blaedel, Acting Center Director

The projects described in the following articles demonstrate the breadth of interest, from machine tool design to manufacturing process engineering, that has traditionally defined the Center for Precision Engineering at Lawrence Livermore National Laboratory (LLNL). In addition, a new set of programmatic requirements has been articulated that will challenge the Center far into the future.

The first article describes a project whose goal is to devise and prove a method of dimensional metrology for non-rigid components, for example, those that can deform during the inspection process. No such methodology exists to the precision that will be required for the weapons program in the future.

The second article summarizes a new approach to designing a precision machine tool where its performance is characterized in the spatial-frequency domain instead of the spatial domain alone, which is the basis for classical design.

The third article describes how we combined data sets from three metrology instruments to achieve a broader characterization of surface quality than could be done by any one of the instruments alone.

The fourth article describes a fabrication process that holds promise for predictably removing material down to the molecular level.

The fifth article discusses a machining process whose goal is to radically reduce the wear of a diamond tool, thus enabling us to machine difficult materials to higher precision than is currently possible.

We will continue to pursue some of these ideas through the next year.

In looking to the future of Precision Engineering, we have recently defined a new challenge that will provide a focus for our R&D resources. This focus is the manufacturing and characterization of "meso-scale devices" for LLNL's National Ignition Facility (NIF). The science-based Stockpile Stewardship Program will conduct weapons physics experiments




in NIF in much the same way it used to conduct experiments at the Nevada Test Site (NTS). The major difference between the NTS experiments and those on NIF will be one of scale. Future weapons physics experiments will be of the order of millimeters in the NIF, which is 100 to 1000 times smaller than those conducted at the NTS. These millimeter-scale physics experiments or "meso-scale devices" will provide data about shock physics, equation of state, opacity, and other essential measurements of weapons physics.

This program faces many new challenges, not the least of which is the ability to design, manufacture, characterize, and field the experiments in the NIF chamber. These experiments will involve millimeter-size structures fabricated from exotic materials and to precise tolerances down to 0.1 μm or less. In addition, cleanliness will be an extremely important requirement since particulate even at the sub-micrometer level could ruin an experiment. Diagnostics must interface with the meso-scale device at precise locations and will include miniaturized fiber optic technology, x-ray optics, digital cameras, and radiation detectors. These experimental systems must be aligned in the NIF chamber to a precision of micrometers over distances of meters. And finally, all of this must be done at minimum cost and expenditure of time since there will be hundreds of experiments performed per year if NIF is to be a cost-effective facility and fulfill its promise of science-based stockpile stewardship.

Virtually all of the engineering and manufacturing problems faced by the early weapons program must be rethought, redefined, and resolved. The problem of small, precise devices, including diagnostics, requires a systems engineering approach to meet the challenges of precision, cost, and throughput. The challenge to engineering in general and to the precision engineering community in particular is clear.

Center for Precision Engineering

3



3. Center for Precision Engineering

Overview

Kenneth L. Blaedel, Acting Center Director

Atmospheric Pressure Plasma Processing for Damage-Free Optics and Surfaces

Jeffrey W. Carr 3-1

Design of an Ultra Precision Machining Facility: ULTRA350

Keith Carlisle and Stanley L. Edson 3-11

Intermittent Single-Point Machining of Brittle Materials

Mark A. Piscotty 3-17

A Spatial-Frequency-Domain Approach to Designing Precision Machine Tools

Debra Krulewich Born 3-21

Power Spectral Density Measurements of Machined Surfaces over an Extended Range of Spatial Frequencies

Jon M. Baldwin 3-25

Dimensional Metrology of Non-Rigid Objects

Kenneth L. Blaedel and David W. Swift 3-29

Atmospheric Pressure Plasma Processing for Damage-Free Optics and Surfaces

Jeffrey W. Carr
Manufacturing and Materials Engineering Division
Mechanical Engineering

The goal of this project is to develop a fabrication technology capable of shaping and finishing difficult materials with minimal subsurface damage. The technique developed uses an atmospheric pressure mixed gas plasma discharge as a sub-aperture polisher of fused silica and single crystal silicon. Workpiece material is removed at the atomic level through reaction with fluorine atoms. These reactive species are produced by a noble gas plasma from trace constituent fluorocarbons or other fluorine containing gases added to the host argon matrix. The products of the reaction are gas phase compounds that flow from the surface of the workpiece, exposing fresh material to the etchant. An initial system has been built and tested. Preliminary etch rates have been measured on fused silica and silicon.

Introduction

Modern materials present a number of formidable challenges to the fabricators of optics and precision components. In many Lawrence Livermore National Laboratory (LLNL) programs there is a general need for spherical and non-spherical surfaces fabricated from a variety of media that are able to handle high fluences. Traditional finishing and final figuring techniques, still largely a black art, fall far short of today's requirements. Unfortunately, any form of contract figuring or polishing introduces a degree of distortion into the workpiece. Present finishing technologies produce significant levels of surface and subsurface damage and are difficult to apply to non-spherical shapes. In many applications, mechanically inferior or more costly materials are used because the material of first choice (in terms of performance) is too difficult to fabricate. Reactive atom plasma polishing presents an opportunity to significantly improve the capabilities of our precision manufacturing technology.

The thrust of this effort is to develop a final procedure that can remove damage introduced by previous process steps, figure highly aspheric optics or precision components and reduce high-frequency surface roughness. In its final stage, this technology functions as a shaping and a finishing operation in both the subtractive and the additive mode. Subsurface damage in the form of cracks and plastically deformed materials are removed in an atmospheric pressure, dry etch

process. Removal rates in an optimized system are expected to be high, and large workpieces will be configured in a relatively short time.

The final shaping of optics and precision surfaces is executed by exposing the workpiece to reactive atoms generated by an atmospheric pressure, inert gas plasma. Reactants are created from trace gases introduced into the discharge and are chosen to fit the chemistry of the workpiece. Interaction of reactants with the workpiece is almost entirely chemical, resulting in the removal of the existing damaged layer without distorting the underlying material. The plasma footprint is highly reproducible, resulting in predictable removal rates. By translating the plasma across the workpiece, a symmetric or highly asymmetric damage-free surface can be created.

Background

In demanding applications the success of any manufactured precision component depends on the surface and near-surface quality. Damage, present in crystalline or amorphous substrates as cracks, voids or plastically deformed material can lower the damage threshold for high fluence use or increase the chemical activity of the surface (corrosion). Any shaping and finishing process that involves surface contact will introduce damage to some degree. In addition, for silica glass optics, high quality final finishing is accomplished through complex solution phase chemistry that tends to deposit contaminants

in a gel-like layer on the surface of the component. While an extremely smooth part can be produced, the top layer often hides an extensive damage zone below. An optimum approach to final surface generation would involve a non-contact system that removes material chemically without residual contamination. A number of attempts at non-contact material removal have been reported.

Plasma etching at reduced pressure is used extensively in the semiconductor industry for processing a wide variety of materials including semiconductors, metals and glasses.¹ Removal mechanisms and removal rates have been studied in detail and are reasonably well understood. With pressures in the 10 to 20 mtorr range, ion and electron densities are on the order of 10^9 to 10^{10} cm⁻³. These reactive ions are believed responsible for the majority of material removal. Consequently, the technique is known as reactive ion etch (RIE). With electron energies in the 3 to 30 eV range, material removal tends to be largely chemical in nature. Below 50 eV, physical sputtering is negligible and subsurface damage is non-existent. Above 50 eV, the induced damage will affect device performance. In the absence of sputtering, reaction products must be volatile or the process is self-terminating after the formation of the first product layer.

Considerable effort has been put into developing plasmas with uniform etch rates over the entire discharge, making RIE unsuitable for the production of figured precision components. Shaped discharges are possible, but usually have a Gaussian shape over the width of the reaction chamber. If removal of the damage layer from a pre-shaped part were the only task, RIE would be usable for smaller optics. However, the technique would fail when non-uniform material removal is required. RIE etch rates are also quite low and not suitable for removing large amounts of material.

An additional practical drawback to the technique for precision finishing would be the requirement of a vacuum. Translating either the source or workpiece with precision in a complicated pattern inside a vacuum chamber would be challenging, especially in the case of large optics. *In-situ* metrology would also be awkward. Fortunately for the development of this process, the chemistry that occurs in RIE systems is similar to that of higher pressure systems. The main difference is the population shift between atom and ion species. In the tail of the high-pressure discharge, the part that interacts with a workpiece, the population of ions would be low. There, atoms produced by ion-electron recombination reactions predominate. However, reactive gas mixtures that are used in RIE will also work for atmospheric pressure

systems. In the atmospheric plasma it is the chemical potential of the atoms rather than the ions that is responsible for the removal of material.

A modified RIE for polishing at reduced pressure has been built using a capacitively coupled discharge.²⁻⁵ Named "Plasma Assisted Chemical Machining" (PACE), the system has been successful in shaping and polishing fused silica. In the first application of PACE, Bollinger *et al.*² noted that a 1-in. discharge controlled material removal to better than 1% with removal rates from 0 to 10 μ m/min. The footprint and the rates could be varied during the process by changing plasma parameters such as power and reactive gas flow. In subsequent analysis, the parts showed no evidence of subsurface damage, no surface contamination and no distortion at the edge of the optic (roll-off). Material removal could be predicted from a linear superposition of the static footprint. Control algorithms were developed to permit fabrication of a spherical optic from knowledge of static behavior.³

An updated low pressure system with a 13-mm capacitively-coupled discharge was built to shape and polish single crystal silicon and silicon over silicon carbide.⁵ While the results for subsurface damage were similar to previous studies²⁻⁴ the authors also considered the evolution of surface roughness and its relationship to previous process steps. Not surprisingly, they found that greater subsurface damage results in an increase in roughness.

A major limitation of the capacitively-coupled discharge is the requirement that the workpiece be either conductive or <10 mm thick. In addition, etch rates were noted to be dependent on part thickness, decreasing by a factor of ten when thickness changed from 2 to 10 mm. Above 10 mm the rates were too low to be of much use for shaping, but could find some application in damage layer removal (20 nm/min).⁴ Optics up to 188 mm in diameter with 30° slope were polished under mild vacuum. If metrology is needed in an iterative procedure, the chamber must be vented, the part removed and measured, replaced in the chamber and then pumped down for the next etch step. The convergence rate for PACE is typically low, resulting in a long and arduous (expensive) process. *In-situ* real-time metrology becomes difficult if the equipment is not dedicated to a single type of workpiece. While extension to larger workpieces is certainly possible, the difficulty in handling large parts is clear.

Ion beam sputtering or neutral ion beam milling remove material from the workpiece by kinetic interaction of ions with atoms or molecules of the surface. The technique has been around for quite a while⁶ and the main application has been optical

polishing for fused silica optics.^{7,8} Earliest sources used beams with energies that were a large fraction of 1 MeV, while most recent systems use Kaufman sources with energies of 1500 eV to provide the optimum sputter yield. Researchers claim several advantages for ion milling. Chief among them are: no surface contact, no weight on the optic, no edge effects, high removal rate, and efficient correction of long spatial wavelength errors. The removal rates for a 1-A beam on fused silica are about 0.35 cm³/h; unfortunately, such high currents are rarely used. A more typical value is 30 mA providing a removal rate nearer to 0.01 cm³/h. Zerodur, a proprietary form of low-expansion glass, was a factor of two slower. Ten hours were required for correcting (not figuring) a 30-cm optic.

As with any sub-aperture tool, the footprint must be stable and predictable. The Kaufman source, which typically produces Gaussian beams 3 to 15 cm in diameter, meets these needs quite well. The size of the beam can usually be reduced with an aperture. Configuring algorithms for the production of aspheric or non-rotationally symmetric workpieces have been extensively developed.⁹

Observed disadvantages include high surface temperatures, an increase in surface roughness, and the need for a vacuum (and to do translation in that vacuum). The temperature is dependent on beam current so an increase in etch rate assumes an increase in temperature, often surpassing several hundred degrees Centigrade. An increase in roughness on undamaged material occurs primarily from redeposition, although later efforts suggest that any amount of subsurface damage in the workpiece will increase roughness and degrade surface quality.¹⁰ Typically, a 50% roughness increase for shallow etched parts can be expected.¹¹ The difficulty in working under vacuum is only a practical consideration, especially for large optics. Clearly this doesn't concern the engineers at Kodak who have built an ion beam figuring chamber capable of handling a workpiece that is 2.5 m × 2.5 m × 0.6 m in size.⁸

One investigator has used a direct current plasma (DCP) at atmospheric pressure to build a device capable of thinning wafers.¹² The system, called a "plasma jet," uses argon as the plasma gas with a trace amount of fluorine or chlorine for reactive atom production. The main intent of the device is to do backside thinning of processed silicon wafers for smart card and other consumer applications. The industry requirement of a 200-mm wafer with a thickness <50 μm cannot be met with any current commercially viable polishing process. The defects and microcracks introduced by abrasive systems create a damage layer that is a large fraction of the desired 50-μm thickness. Thin wafers produced by

polishing are prone to fracture even with delicate handling. Wafers are thinned in a batch mode by placing them on a platen and using planetary type motion to move the sub-aperture plasma in a pseudo-random fashion across the surface. Irregularity from fluctuation in the plasma jet is removed by this randomization. The discharge is about 1 in. in diameter and the removal rate is 0 to 20 μm/min for a 200-mm wafer with a uniformity of <5%. Total material removal comes to about 30 cm³/hr.

A number of analytical methods were used to assess surface quality. Scanning electron microscopy did not show defects or scratches on either side of the wafer. X-ray photoelectron spectroscopy, used to measure surface contamination, showed no evidence of elements other than silicon. The chemistry of the plasma is quite specific. As a result, metal or carbon contamination (fingerprints) present before etching were not removed. Transmission electron microscopy did not reveal any subsurface defects in the silicon crystal, supporting the supposition that the plasma is nearly 100% chemical in nature. A number of other studies, including optical microscopy and adhesion tests, did not reveal any additional defects.

In its current configuration, the DC plasma is not well suited for aspheric generation or material deposition. The trace reactants are introduced along with the bulk gas and, as a consequence, are widely distributed throughout the discharge, substantially increasing the footprint and the minimum feature size. Electrodes that are used to establish the arc are eroded by the reactants and add particulates to the atom stream. Electrode erosion also causes fluctuations in the plasma and accounts for the reduced uniformity compared to RIE systems. Detrimental electrode reactions also preclude the use of oxygen and many other potentially useful plasma gases. Finally, the discharge is not as hot as an inductively-coupled plasma (ICP) and, as a result, the production of reactant atoms is reduced.

A radio frequency (RF) plasma has been used to slice silicon and as a sub-aperture tool to polish optics.¹³⁻¹⁵ The plasma is generated around a wire or blade electrode immersed in a noble gas atmosphere that contains a trace of reactive components. The plasma converts the reactive precursors to radical atoms that react chemically with the workpiece, removing material one atom at a time. Referred to as chemical vapor machining (CVM), the electrode is brought to within 200 μm of the workpiece wherever material is to be eroded. Analysis has shown the resulting surfaces to be damage free and the process is considered to be entirely chemical in nature. A comparison of damage in silicon for

polishing, sputtering, CVM and chemical etching is reported in the literature.¹³ Both mechanical polishing and argon sputtering induced significant damage into the silicon surface. The damage for CVM and wet chemical etching was similar and near to the intrinsic damage typically found in silicon used in the semiconductor industry.

Unfortunately, the non-rotationally symmetric nature of the CVM footprint makes the process difficult to model and control. In the case of a blade, the footprint takes the shape of a high aspect ratio rectangle with rounded corners. Process rates are limited by the ability of the plasma to decompose reactive precursor into radical atoms. While no vacuum is required, the workpiece must be enclosed in a vessel to contain the plasma atmosphere.

The ICP discharge in its most familiar commercial form was originally developed to grow crystals.^{16,17} In a configuration remarkably similar to the excitation source used in current analytical spectrometers, a powder of the crystal to be synthesized was aspirated into the center of the discharge. Reed was able to grow boules from 5 to 15 mm in diameter and as long as 30 to 90 mm with a growth rate of 20 to 50 mm/hr. No mention was made of crystal quality. This approach to crystal growth seemed dormant until the early part of this decade when it was used to produce crystalline films of a number of oxides (MgO , ZrO_2 , NiO , SnO_2 , TiO_2 , ZnCr_2O_4 , Cr_2O_3 , CoCr_2O_4 , NiCr_2O_4 , and several rare earth oxides).^{18,19} X-ray diffraction was used to confirm the crystalline nature of the films. Superconducting thin films of Bi-Pb-Sr-Ca-Cu-O were also fabricated with plasma spray methods.²⁰ Between the two efforts, an ICP was used on several occasions to produce ultrafine particles by desolvating droplets aspirated into the discharge.²¹

The conditions of the ICP make it an ideal source for reactive atoms needed for shaping damage-free surfaces. In his initial work Reed surmised that the high electrical conductivity of partially ionized gases (for argon, $120 \text{ } \Omega/\text{cm}$ at $15,000 \text{ K}$) contributed to the ease of plasma formation at high pressures. There are no electrodes and a number of gases can be used as the host matrix, although argon is usually the principle component. A typical discharge is characterized by high current (100 to 1000 A) and a relatively low voltage (10 to 100 V). The flowing plasma is not in complete thermodynamic equilibrium; however, ion and excited state atom populations are within 10% of equilibrium values. Electron densities are high ($>10^{15} \text{ cm}^{-3}$ is typical) suggesting equilibrium temperatures $>9000 \text{ K}$.²²⁻²⁴ Reed calculated a peak temperature of $10,000 \text{ K}$ from the ratio of emission intensities for a set of argon lines, again assuming equilibrium,

making the ICP an efficient source for the generation reactive atoms.

Figure 1 shows a standard plasma torch used as an emission source for analytical spectroscopy. The current from a 27.12-MHz RF generator flows through a three-turn copper load coil around the top of the torch. The energy is coupled into the plasma through a cylindrical skin region that is located on the outer edge of the discharge nearest the load coil. The plasma is supported in a quartz tube by the plasma gas that is introduced tangentially to form a stabilizing vortex. In analytical spectroscopy, trace elements are introduced into the plasma as an aerosol through the central tube. The plasma skin is thin in this region and the sample easily penetrates the discharge. As the sample travels upward it




Figure 1. ICP torch showing the three types of gas flow and the location of the plasma relative to the load coil.

becomes progressively desolvated, atomized, excited and ionized. The relative distribution of ions and atoms in the discharge is represented in **Fig. 2**. Spatial profiles for five plasma regions indicate that the ion population decays before that of the excited atoms. The maximum atomic emission occurs several millimeters above the load coil near the visible tip of the discharge (zones 3 and 4). Recombination and radiative decay in this region are used to spectroscopically determine the concentration of the sample in the injected solution. Since the main thrust of ICP research has revolved around its use as an excitation source, information on the unexcited neutral atom population throughout the discharge is limited.

The characteristics of several shaping and finishing technologies are summarized in **Table 1**. Typically, all of the approaches have one or more areas where they excel. In the case of low-end consumer optics where cost is always an issue, traditional methods such as lapping and polishing will be used for many years. For some applications, including high fluence optics, the categories of subsurface damage, contamination and figure control would be emphasized. Unfortunately, none of the current techniques is adequate for demanding applications. It is in these critical areas that reactive atom plasma polishing (RAPP) will be very effective. No other method has the potential to produce damage-free, non-contaminated surfaces with high removal rates and good figure control over a large range of materials.

Progress

An initial system has been designed and built using commercial equipment as well as home-built components. The unusual application of the RF generator and impedance-matching network made initial operation difficult. Several modifications to the load coils and plasma chamber were necessary to achieve a stable discharge. The range of gas flow rates necessary to achieve a stable plasma was studied. Temperature shifts at the workpiece surface were measured over a period designed to simulate the figuring of a highly aspheric optic. Etch rates were measured for silicon.

Design and Construction

A commercial plasma generator from Advanced Energy, developed for use as a low-pressure etching system for use in the semiconductor industry, has been modified to produce reactive atoms in a flowing argon stream. The plasma chamber, containing the torch and load coils (**Fig. 1**) was built from copper and attached directly to the impedance-matching network. Initially, a four-turn coil was used to couple energy into the discharge. The energy from the RF generator was shunted directly through the water-cooled coil, constructed from 3-mm copper tubing. Unless the plasma is ignited, there is very little resistance to the energy flow. The 5 kW from the generator either reflects back to the source or flows to ground. After ignition the power is inductively

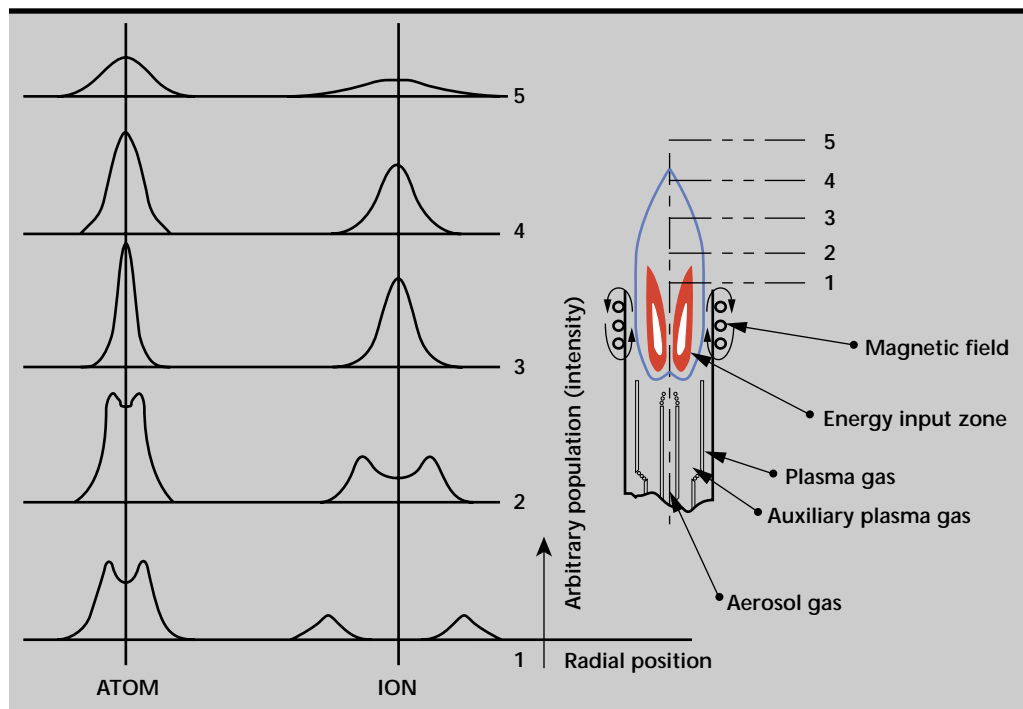


Figure 2. Atom and ion distribution in the ICP.

Table 1. Comparison of finishing and shaping methods.

	Material set	Figure control	Removal rates	Sub-surface damage	Surface roughness	Contamination	Foot-print	Vacuum required
Diamond machining	Moderate	Very good	Very high	High	Good	None	1 to 100 μ m	No
Lapping and grinding	Large	Very good	Moderate to high	Very high	Moderate	Moderate	25 mm to full aperture	No
Chemical etch	Large	None	High	None	Good	Low	Full aperture	No
Traditional polishing	Small	Poor	Low	High	Excellent	High	25 mm to full aperture	No
RIE	Moderate	Good	Moderate	Low	Good	Low	Full aperture	Yes
PACE	Small	Good to very good	Moderate	Low	Good to very good	Low	25 mm to full aperture	Yes
Ion milling	Large	Good	Moderate	Moderate	Good	Low to moderate	100 nm to 50 mm	Yes
Plasma jet	Small	Good	High	None	Good	Possible	10 mm to 25 mm	No
Plasma CVM	Moderate to large	Moderate	Moderate	None	Good to very good	None	Long and narrow	No
RAPP	Moderate to large	Good to very good	Very high	None	Good to very good	None	1 mm to 50 mm	No

coupled into the plasma; very little power is either reflected or lost to ground. Under normal operating condition, the discharge generates large amounts of stray RF energy that must be confined to the plasma chamber. The design must take these two operating extremes into account.

The plasma chamber was constructed from a single sheet of 0.125-in. copper, folded into the proper shape. The use of folds rather than seams reduces the amount of RF leakage that invariably occurs at the joined edges. Where connection was necessary, the bond was created by braising the pieces together, minimizing RF leakage. An inside view of the copper box (**Fig. 3**) shows a later model torch mounted in place of the earlier solid version. A

two-turn coil also replaced the original four-turn design. The lower copper block on the rear wall is the hot end of the coil and is isolated from the plasma chamber. The ground side can clearly be seen, firmly attached to the inside of the box. In the configuration pictured here, a long outer torch tube has been added to the torch to prevent quenching of active species via entrained atmospheric gasses. The torch extends into the sample chamber (**Fig. 4**).

The discharge occurs in the region of the coils (**Fig. 5**). The shadows of the two turns of the load coil can be seen through the netting of the observation window. They are located just to the left of the bright portion of the plasma. The hot gases are contained entirely within the quartz outer tube of the

torch. The color of the plasma is modified by the trace elements used to form the reactive atoms. In this photograph, only the argon host gas is being excited. Nearly all of the energy (1.5 kW) used to




Figure 3. View inside the plasma chamber showing the demountable torch with a long outer tube.




Figure 4. The sample chamber. The three-axis translation stage can be seen through the observation window. The plasma chamber is to the left.



Figure 5. The discharge seen through a dark filter glass.

generate this discharge is radiated in the visible and near ultraviolet region of the spectrum. The light is produced through relaxation of excited states of argon as well as ion/electron recombination of the ionized host gas. Only a small portion of the energy is diverted to kinetic energy and to dissociation of the reactive atom precursors.

Reactive atoms from the plasma are delivered to a workpiece that is mounted on translation stages in the sample chamber. In its present configuration, a crossed X-Y slide and a rotary stage allow the fabrication of rotationally symmetric workpieces up to 200 mm in diameter. Repositioning both the torch and the sample mount within their respective chambers can increase the range of the stages. A large range was necessary to investigate conditions for material deposition in the additive configuration of the reactive atom system. A large ventilation duct removes reaction products and the host argon gas along with any unreacted chemicals from the sample chamber.

Initial Operation and System Modifications

The plasma was initially tuned on argon gas only. Before any trace gases were added, the plasma was stable over a wide range of tuning and operating conditions. The auto tuning function supplied with the system for low-pressure operation was able to compensate for any shifts in impedance when power settings or gas flows were altered. The system was tuned to operate under minimum reflected power. When run at 1.5 kW, a typical operating condition, the reflected power was below 15 W.

The first gas used to generate fluorine atoms was a mixture of sulfur hexafluoride (SF_6) in a nitrogen host. Unfortunately, the nitrogen seemed to promote quenching of the reactive atoms, lowering the etch rate to nearly zero and complicating plasma tuning. The precursor was switched to pure SF_6 reducing quenching and improving tuning sensitivity. After a prolonged period of use, SF_6 appeared to deposit sulfur on the surface of the plasma chamber as well as the sample box and the ventilation system. Energy in the radio frequency range tends to propagate along the surface of a metal. Adding a dielectric to the surface in the form of a sulfur compound substantially altered the system tuning characteristics.

To adjust to the new tuning conditions, it was necessary to reduce the number of turns in the load coil and to shift the spacing between the coils. Even after a considerable investment of time, the system never regained its wide stability range. The auto tune function is now unable to compensate for shifts during operations. It is now necessary to tune

manually for minimum reflected power. When a suitable tuning point is found, reflected power again stays below 15 W; however, the adjustment range of the tuning capacitors is very narrow (that is, a narrow range of stability).


During early operation of the reactive atom plasma, the exterior of the plasma chamber and sample box was measured for RF leakage. The detection system was calibrated by placing it inside the plasma box without igniting the plasma. The generator was then operated over a range of power from 0.1 to 1.0 kW to estimate the sensitivity of the detector. We were then able to measure power levels outside the plasma chamber and sample box as well as inside the sample box over the entire range of operating conditions with and without plasma ignition. Not surprisingly, radiation was detected in the workpiece area, although the energy levels were very low when the plasma was ignited. Outside the plasma chamber and sample box, it was not possible to detect any RF signal.

To limit reactive atom quenching, the length of the outer torch tube was increased in several steps until it nearly touched the workpiece. For this experiment, a demountable torch was used. The demountable system permits a change in the tube length without dismantling of the entire torch assembly. It is also useful on those rare occasions when something goes amiss and the outer tube is transformed into a glowing pile of quartz slag. While reaction rates at the surface of the workpiece did increase, suggesting an increase in the reactive atom population, surface temperatures also increased. Thermocouples were placed at several locations in the sample box; one above the sample, one between the thin silicon wafer sample and the rotary stage and two others at probable hot spots in the chamber. The maximum temperature of the workpiece (measured from the rear of a silicon wafer) increased from 70 °C with the short torch to 225 °C with the longest torch. The temperature is reached quickly, a matter of a few minutes, and stays there for the duration of the experiment, as long as 60 min in some cases. At this time, the major concern with the temperature increase is resistance of the stage motor. To prevent overheating, the rotary stage and motor were covered with a cooling coil.

Silicon Etch Rates

A silicon wafer with a known mass was exposed to a flow of fluorine atoms from the reactive atom plasma. The fluorine atoms were created from the dissociation of SF₆ delivered to the plasma at a

flow rate of 30 ml/min. In 20 min the silicon wafer was reduced in mass by 0.372 g. In a balanced reaction, two molecules of SF₆ react with three atoms of silicon to produce three molecules of SiF₄ and two atoms of silicon:



If the reaction were 100% efficient a total of 1.143 g of silicon would have been consumed. Under the operating parameters used in this experiment, silicon reacted with an efficiency rate of 32.6%. The reduction in rate could be from a number of factors. The two most probable are quenching or the inability of the reactants to reach the surface of the workpiece. Quenching of the reactive gases in the plasma can come from at least two sources: either from entrained atmospheric nitrogen or from recombination of the fragmented precursor. At the surface of the workpiece, boundary layer conditions may prevent the fluorine atoms from reacting. Two other experiments were performed under similar conditions, the only change being small changes in flow rate of the reactive gas and a small increase in the plasma power.

Initial studies using an optical interferometer have shown the footprint to be roughly Gaussian in shape. The material is removed in a symmetric fashion across the discharge. As expected, there is little evidence of redeposition. In addition, preferential etching, caused by dirt or other debris on the surface seems not to be an issue in early tests.

Future Work


In the initial experiments, the reactive precursor was added to the plasma, giving the reactive atoms the broadest distribution possible. The concentration of reactive species was also kept on the low side. The combination of a wide footprint and a low reactant concentration should minimize the boundary layer effect. On the surface, the 30% efficiency suggests that quenching in the plasma was responsible for a reduction in active species leading to the lower than expected consumption of silicon.

To confirm this statement, the proficiency of the reaction over a wide range of concentration will be measured. The distribution of reactants throughout the discharge will also be altered. Confining the reactants to the central channel of the plasma will decrease quenching of the active species. The addition of a second downstream load coil can also serve to redissociate the fluorine if recombination is an issue. Besides increasing understanding of the

reduction in reaction rate, this study will determine the range over which the reaction is linear, an important parameter in fabrication of precise optical profiles. The etching pattern caused by a translated discharge will also be studied.

The shape and symmetry of the static footprint will be measured as a function of plasma conditions (power, shape of load coils), primary gas mixture, primary gas flow rate, reactive gas type and concentration, discharge size, distance from the plasma and dwell time. The effect of crystal orientation on the etch rate of silicon will also be investigated. The footprint shape, surface roughness and removal rates will be measured with a Fizeau-type interferometer and an atomic force microscope. If necessary, etch rates and materials removal mechanisms will be determined *in situ* by emission or absorption spectroscopy. Once the stability and shape of the footprint is determined, convolution algorithms can be developed to generate specific surface shapes. The program would determine a path of motion for the discharge over the surface of the workpiece. A long term objective is to develop a well understood technology that can be tuned to produce a highly asymmetric surface in a single operation from a data file.

References

1. G. S. Oehrlein (1994), "Reactive ion etching," *Handbook of Plasma Processing Technology*, Vol. **196**.
2. L. D. Bollinger and C. B. Zarowin (1988), Part 1, *Proc. SPIE*, Vol. **966**, p. 82.
3. L. D. Bollinger and C. B. Zarowin (1988), Part 2, *Proc. SPIE*, Vol. **966**, p. 91.
4. S. J. Hoskins and B. Scott (1995), *Proc. SPIE*, Vol. **2542**, p. 220.
5. S. J. Hoskins and B. Scott (1995), *Proc. SPIE*, Vol. **2542**, p. 235.
6. A. Meinel, S. Bashkin, and D. Loomis (1965), *Appl. Opt.*, Vol. **4**, p. 1674.
7. S. R. Wilson, D. W. Reicher, and J. R. McNeil (1988), *SPIE*, Vol. **966**, p. 74.
8. L. N. Allen and R. E. Keim (1989), *SPIE*, Vol. **1168**, p. 33.
9. T. W. Drueding, T. G. Bifano, and S. C. Fawcett (1995), *Prec. Eng.*, Vol. **17**, p. 10.
10. S. R. Wilson and J. R. McNeil (1987), *SPIE*, Vol. **818**, p. 320.
11. S. C. Fawcett, T. W. Drueding, and T. G. Bifano (1994), *Opt. Eng.*, Vol. **33**, p. 967.
12. O. Siniaguine (1996), *Int. Symp. Plasma Process-Induced Damage*, Vol. **151**.
13. Y. Mori, K. Yamamura, K. Yamauchi, K. Yoshii, T. Kataoka, K. Endo, K. Inagaki, and H. Kakiuchi (1993), *Nanotechnology*, Vol. **4**, p. 225.
14. Y. Mori, K. Yamamura, K. Yamauchi, and Y. Sano (1996), *The Japan - China Bilateral Symp. on Adv. Manuf. Eng.*, Hayama, Kanagawa, Japan.
15. Y. Mori, K. Yamamura, K. Yamauchi, K. Yoshii, T. Kataoka, K. Endo, K. Inagaki, and H. Kakiuchi (1993), *Technology Reports of Osaka University*, Vol. **43**, No. 2156, pp. 261-266.
16. T. B. Reed (1961), *J. Appl. Phys.*, Vol. **32**, p. 821.
17. T. B. Reed (1961), *J. Appl. Phys.*, Vol. **32**, p. 2534.
18. M. Suzuki, M. Kagawa, Y. Syono, and T. Hirai (1990), *J. Crystal Growth*, Vol. **99**, p. 611.
19. M. Suzuki, M. Kagawa, Y. Syono, and T. Hirai (1991), *J. Crystal Growth*, Vol. **112**, p. 621.
20. M. Suzuki, M. Kagawa, Y. Syono, T. Hirai, and K. Watanbe (1991), *J. Mat. Sci.*, Vol. **26**, p. 5929.
21. M. Kagawa, M. Kikuchi, R. Ohno, and T. Nagae (1981), *J. Am. Ceram. Soc.*, Vol. **64**, p. C7.
22. "Inductively Coupled Plasma Emission Spectroscopy, Part I: Methodology, Instrumentation and Performance, in Chemical Analysis" (1987), Vol. **90**, P. W. J. Boumans, ed., John Wiley & Sons, New York, New York.
23. "Inductively Coupled Plasma Emission Spectroscopy, Part II: Applications and Fundamentals in Chemical Analysis" (1987), Vol. **90**, P. W. J. Boumans, ed., John Wiley & Sons, New York, New York.
24. *ASPE Spring Topical Meeting on Prec. Eng.* (1998), Carmel, California, April. 

Design of an Ultra Precision Machining Facility: ULTRA 350

Keith Carlisle and Stanley L. Edson
*Manufacturing and Materials Engineering Division
 Mechanical Engineering*

This report describes the advanced machine tool technology used in the design of an Ultra Precision Machining Facility, the ULTRA 350 at Lawrence Livermore National Laboratory (LLNL). It outlines the collaborative research undertaken in key technology areas to ensure we maintain our world leadership in precision machine technology. It also describes the new enhanced manufacturing process capabilities the machine brings to LLNL's precision manufacturing group.

Introduction

The drive for miniaturization and higher precision across a widening range of manufacturing activities is increasing, placing demands on engineering for new capabilities for machining, positional control, and measurement to nanometer tolerances. It also calls for improved understanding of materials and processes at the molecular or atomic level.

NASA recently reported that one of the principal missions for the International Space Station (ISS) is material science in microgravity, stating that "manufacturing is 17% (\$1.2 trillion) of the US gross domestic product." This means that even the modest improvements in materials and their production can have a great impact on the US economy. The 1998 Metal Casting Industry Roadmap lists lack of knowledge of interactions among process, microstructure, chemistry, and material properties as one of the major technology barriers in materials.

It is critical to the ongoing success of LLNL and the US economy that we have the right engineering tools that will lead to improvements in our manufacturing techniques and the development of new products and processes. Our objective is to establish a national world-class research and manufacturing facility to take us confidently into the twenty-first century.

Progress

The ULTRA 350 machine concept design is now complete and the final detail design will commence in the next fiscal year.

The machine configuration is illustrated in **Figs. 1 and 2**. **Figure 1** shows the machine as it would be used in a single-point diamond turning mode and **Fig. 2** shows the machine in its grinding mode.

This four-axis diamond turning and grinding machine uses a 'L'-bed base construction, which supports the 'X' axis tool carriage slide and the 'Z' axis work slide. Both carriages run on pressurized hydrostatic oil guide-ways. The 'Z' work carriage carries the 'C' axis work-head spindle, which has hydrostatic bearings. The 'X' tool carriage carries the integral 'A' axis rotary table which also has hydrostatic bearings. This rotary table supports a tool-post for diamond turning or metrology equipment for workpiece inspection. A good feature of




Figure 1. Schematic illustration of ULTRA 350 in single-point diamond turning mode.




Figure 2. Schematic illustration of ULTRA 350 in grinding mode.

this rotary table is that it can be removed from the machine to accommodate a grinding spindle.

The work capacity of the machine is 350 mm in diameter and 225 mm in length. Each slide has a maximum travel of 350 mm with a positional resolution of 1 nm. The target performance for the machine is 25 nm (p-v) contouring accuracy and 4 nm rms surface finish.

The machine base is to be constructed from steel, chosen because of its high modulus of elasticity, giving high specific stiffness. Finite element analysis will be used to determine the final design, which may include the selective addition of synthetic granite to achieve high internal damping. Subjecting the base fabrication to thermal stress relief will ensure long-term dimensional stability. Seismic vibrations will be reduced by supporting the base on a three-point, self-leveling pneumatic isolation system, which may have to stand on a damped, large-mass concrete foundation. This will depend on the final site conditions.

Oil hydrostatic bearings have been selected for their attainable high stiffness and load carrying capacity. They also provide smooth motion and good damping with freedom from slip-stick, all essential requirements when servo-positioning slides down to 1 nm during contouring operations. Oil also has the added advantage of being temperature controllable to within 0.01 °C or better to ensure thermal stability of the machine.

The work slide uses an opposed bearing pad design for both vertical and horizontal bearings. Designing the slide motion to be parallel to the work spindle axis ensures that any out-of-balance forces generated by the rotary motion will be absorbed by the slide bearings and not seen directly by the linear drive system.

The tool slide uses the same bearing design as the work slide. Mounting the bearings to a vertical face gives better operational access. It also enables the 'A' axis rotary table to be embedded within the carriage and permits greater slide bearing separation, which was an important design consideration when trying to reduce carriage roll effects. Carriage roll is seen as a departure from straightness at tool travel height in the 'Z' direction. With a 2:1 ratio between bearing separation and bearing-to-tool position, it is possible to control straightness of carriage motion through bearing adjustment. Another design feature of the machine is the ability to exchange the rotary table with a motorized grinding spindle.

The work spindle is integrated into the work carriage for maximum stiffness and uses oil hydrostatic bearings designed to give high stiffness, high load carrying capacity, good damping and low error motion.

Thermal stability is to be maintained by careful selection of the materials used for its construction: the bearing is designed for low Δt and temperature control of the oil to better than ± 0.001 °C. All these design factors have been optimized to achieve the highest standards in single-point turning and enable full exploitation of ductile regime grinding of brittle materials.

The spindle is to be driven through a non-influencing drive coupling by a high-performance brushless DC motor, supported by an air bearing. The drive motor has its own temperature control system. A rotary encoder is built into the spindle to provide commutation for the servodrive and positional information for a "fast-tool-servo" (FTS) axis.

Additional facilities built into the spindle include provision for two-plane dynamic balancing and vacuum work holding.

The 'B' axis rotary table gives great flexibility to the range of work that can be carried out on the machine. Unlike most other machines, there is no loss in workpiece swing diameter when using the rotary table. This is because of the unique design of the rotary table, which is embedded within the tool slide. The rotary table uses the same hydrostatic bearing design as the work spindle to ensure high stiffness and load carrying capacity. The drive system consists of a direct drive frameless DC torque motor and tacho-generator mounted directly onto the spindle. The measuring system comprises a rotary encoder. The whole assembly is temperature controlled by the oil to better than ± 0.001 °C.

Brushless linear drive motors are used to drive both carriages. The particular non-contact linear servomotor selected for this application possesses

clear advantages over conventional drive train technologies in applications requiring high resolution and high-accuracy positioning. Because linear motors consist of only two non-contacting parts, a moving coil and rare earth magnet track, they eliminate backlash, windup, wear and maintenance issues associated with ball screws. They can also outperform the tried and tested capstan drive system.

The moving coil assembly of this motor design uses a compact, reinforced ceramic epoxy structure without metal support for the coils in the magnetic field. This feature eliminates eddy current losses, cogging and magnetic attraction.

The unique magnet spacing in the magnet track minimizes force ripple and allows for excellent velocity control. This highly rigid, low inertia assembly achieves outstanding motion smoothness by sinusoidal commutation, with or without Hall Effects sensors. The very highest positional resolution and accuracies are possible.

A Heidenhain incremental linear encoder, with measuring steps of 1 nm, was selected to measure linear displacement for each axis. Collaboration with Heidenhain has already led to the development of a new encoder system that can measure in two orthogonal directions so that linear displacement and departure from straightness of carriage motion can be determined simultaneously. This is now the subject of an LLNL patent application.

The machine uses two of these linear encoder systems, which are mounted orthogonal to each other on a super invar metrology frame. Each system uses three reading heads.

One head measures displacement in the direction of the carriage travel and the other two heads measure the orthogonal effects of carriage roll and yaw to determine straightness and tilt. Care has been taken to position the scales at the operational height of the machine to avoid pitch errors. From these measurements, an actual position of tool-to-workpiece can be determined.

Particular attention has been given to the work head where, in addition to the design features described, capacitance sensors are positioned within the metrology frame measuring loop-to-monitor any axial or tilt movement of the face plate. These measurements are to be used for automatic compensation made through the machine computer numerical control (CNC) system.

The optical tool setting facility is a non-contacting device used to measure the tool profile and its relative position to the machine coordinate system to an accuracy of 25 nm. This device enables tool errors to be eliminated from the part profile automatically through the CNC system, resulting in higher accuracy of the part profile.

The FTS device is an additional servo-controlled axis mounted to the tool slide carriage to provide tool motion that is synchronized to the work spindle rotation. A piezoelectric (PZT) mechanism provides tool motion up to a 500- μm maximum travel with a bandwidth of 600 Hz. The PZT is driven by the machine digital signal processor (DSP)-based controller, synchronized to the spindle by an encoder, which provides a master time based in an electronic cam configuration.

This device will be used to improve the intrinsic accuracy of the work spindle, which should be of the order of 25 to 50 nm. By recording the systematic errors of the work spindle in a lookup table, it is possible to servo the FTS in a controlled manner to achieve true nanometric accuracy.

The FTS will also be used to enhance the machine manufacturing capability by enabling non-rotational symmetric components such as correcting secondary mirrors, anamorphic optics, off-axis conic sections, turned on-axes and generated surfaces.

The grinding facility for this machine has been designed for maximum loop stiffness. Unlike many other machines, where the grinding spindle is treated as an attachment that can be somewhat limited in dynamic performance by its support system, the objective for this machine is to design a fully integrated facility.

The air-bearing grinding spindle is designed to use existing Westwind spindle bearing components, and will be driven by a direct drive three-phase motor, which is temperature controlled. The system is mounted rigidly to the tool slide, as can be seen in **Fig. 3**.




Figure 3. Schematic illustration of ULTRA 350, with air-bearing grinding system and tool slide.

Great importance is given to the method of mounting the spindle to achieve a stiff and well-damped support system. The flange mounting interface and ribbed column support to the wheel spindle ensure this.

An essential feature of any grinding facility is the ability to true the wheel to run concentric to the spindle axis. It is also necessary to be able to contour the wheel profile to give clearance to the wheel when grinding concave shapes. The on-machine truing facility uses an air-bearing spindle, which is attached to a servo-driven vertical 'Y' axis carriage. Wheel profiles are to be programmable through the machine CNC system.

Thermal stability of the machine is essential for maintaining geometric and dimensional accuracy.

Hydrostatic oil from the linear and rotary bearings are to be temperature controlled to within ± 0.01 °C accuracy using a control resolution of ± 0.001 °C. Temperature sensors are designed into the machine to monitor both the oil feed and exhaust temperatures to ensure thermal stability of the machine. After exhausting from the machine bearings, the oil is to be deliberately run over the machine surfaces before running back to the collection tank.

All drive motors are to be temperature controlled to the same level. Unlike the bearing oil, the cooling oil from these motors does not come into contact with the machine structure but will be taken back to the collection tank via lagged pipes.

A process coolant facility will be provided for both single-point machining and grinding. The coolant is prevented from contacting the machine structure to avoid thermal transients that may be present due to process heating or chilling effects caused by latent heat of evaporation of the coolant.

Self-supporting machine containment will be used to protect the machine from environmental effects. This also protects the machine from inadvertent collisions or loading from the operator. Acoustic damping material may be added to the containment panels to reduce airborne effects that may induce structural vibration of the machine elements.

The containment also supports an oil shower thermal control system for the machine structure. A laminar flow of oil, controlled to ± 0.01 °C, will be directed over the structure in a manner that will provide good operator working conditions. The oil will be gravity collected in a tray beneath the machine, thus providing insulation from any ground thermal effects.

A Delta Tau, Programmable Multi-Axis Controller (PMAC) CNC controller system has been selected for the machine. Using an open architecture design with standardized PC components, this unit combines a network-ready Pentium-based PC, a multi-axis motion control board and Windows-based operator software.

At the heart of the system is the Turbo PMAC motion control board, which is a high-performance Motorola DSP chip. The feature offers high-speed computational capability for more axes of control, more sophisticated control on each axis, or both. The additional axes can be used as true physical axes or as "virtual axes" for inverse kinematics, superimposed moves, or hybrid force/position loops for grinding operations. The Turbo also offers a special dynamic look ahead/look back buffer that permits acceleration control anticipated over 100 blocks ahead, and reversal through hundreds or even thousands of blocks. This is very desirable for obtaining high accuracy of complex shapes especially when using the FTS axis for non-axis symmetrical shapes.

Future Work

The ULTRA 350 is targeted both for continuing multi-program support and also for enabling program development.

First, this machine will allow us to improve support to our current customers. For example, because of its FTS and increased accuracy, it will allow us to fabricate meso-scale device components that we currently cannot produce. One case is sine-wave interfaces where the pattern of the sinusoidal surface requires a fast-moving axis. In the short term, this machine will provide this improved capability and will more generally allow physicists wider choice of experimental configurations for fielding on LLNL's Omega and NIF projects. In the long term, the ULTRA 350 will lead to the specific upgrade of capability in manufacturing high-precision devices.

Second, this machine will lead to new program development. For example, it will allow certain classes of optical surfaces to be turned that cannot now be turned. One category is an off-axis conic surface where the distance that the surface falls from its optical axis is greater than the swing of the machine. A second category is a surface that is both reflective and diffractive, particularly on an aspheric surface with changing blaze angle. This has value to





Figure 4. Illustration from NASA's Next Generation Space telescope.

internal programs dealing with novel optical designs as well as external agencies such as NASA which require a 6-m-aperture telescope for their Next Generation Space Telescope (NGST) (**Fig. 4**). LLNL has put forward a manufacturing proposal to

machine the 2.1-m-long segments for the telescope, which will require new machining technology to be developed, the third category of anamorphic surface, made possible by the FTS.

Third, this machine will become the platform for the further development of precision machining processes. One example is the machining of brittle materials to complex shapes. Using the highest quality platform is a requirement for such projects to provide the highest probability of success. The Air Force has a continuing work program with LLNL to manufacture optics for the Space Based Laser Program.

Fourth, this machine will evolve to an integral part of our continuing mission to produce the larger weapons components because of its integrated rotary tool axis.

Finally, this machine is a vehicle that allows us to import external, newly developed technologies to LLNL. In a related role, we have assembled a team in anticipation of future projects for programs, such as the upgrade of our facility to manufacture meso-scale devices for NIF. 



Intermittent Single-Point Machining of Brittle Materials

Mark A. Piscotty
*Manufacturing and Materials Engineering Division
 Mechanical Engineering*

Single-point diamond turning is generally considered ill-suited to machine brittle materials, largely due to excessive tool wear resulting in poor part quality. The goal of this project is to investigate observed reduced tool wear by applying a controllable intermittent cutting duty cycle to the diamond turning process of brittle materials. Two methods are evaluated where the cutting edge of the tool is periodically separated from the part surface. It is hypothesized that the tool edge may wear less due to resulting lower temperatures of the diamond. The first method uses piezoelectric (PZT) ceramic elements to induce an elliptical motion of the diamond tool at frequencies ranging from 2 to 30 kHz. The elliptical motion of the tool is used to enhance a typical cylindrical turning operation. A second method investigated is diamond tool fly cutting of flat surfaces also using the concept of intermittent machining, but at a much lower frequency than the ultrasonically assisted turning. Repeatable, low tool wear in currently non-diamond turnable materials holds significant benefits to national defense programs and industry.

Introduction

The demand for innovative fabrication methods compatible with brittle materials is largely driven by organizations such as Lawrence Livermore National Laboratory (LLNL)'s Weapons and Lasers Programs, DOD, and private industry. Single-point diamond machining is an attractive process due to the vast experience and knowledge base available, combined with process efficiency and accuracy. However, these attractive characteristics are typically present only when machining ductile materials such as aluminum, copper or electroless nickel. The goal of this project is to extend the database of important engineering materials that may be diamond turned in an efficient and repeatable manner.

The approach taken here is to vary the cutting dynamics of the material removal process by varying the interaction between the diamond tool edge and the workpiece material. One method of accomplishing this is to vibrate the diamond tool so that it moves in an elliptical motion in the direction of cut, as shown in **Fig. 1**. The amplitude and frequency of the vibration are two parameters that are used to modify the cutting process.

Another method of introducing dynamics modification to the cutting process is to arrange the workpiece and tool in a fly cutting arrangement. In this case the workpiece is typically stationary and the diamond tool moves in a circular planar path relative to the face of the workpiece. The material removal process in this arrangement is intermittent with the on/off contact time determined by the angular speed of the tool, the relative sizes of the workpiece and the circular path of the tool.

Both of these techniques offer potential to modify the standard diamond turning process dynamics to accommodate traditionally non-turnable materials




Figure 1. Schematic of elliptical vibration-assisted diamond turning.

such as brittle materials¹⁻² and steels.³ A leading hypothesis generated to explain the observed reduction of tool wear using these techniques centers about the thermal interaction between the diamond tool edge and the workpiece surface. Specifically, it is hypothesized that catastrophic tool wear can be mitigated by modifying the material removal process such that the diamond tool edge temperature is maintained below a critical temperature at which edge breakdown occurs.

The two methods described above are used in this project in an effort to verify this hypothesis and to develop processes needed to single-point diamond machine materials outside the current range of those used in typical single-point diamond turning.

Elliptical Vibration-Assisted Diamond Turning

A method of PZT-induced vibration at the interface between the diamond tool edge and workpiece surface was developed at LLNL, leveraging both internal ultrasonic and diamond turning expertise, as well as external, international research efforts.^{2,4} The hardware developed for this project includes a stainless steel 'resonator' that holds the diamond tool for the turning application. **Figure 2** shows a photograph of the actual resonator hardware. The resonator is driven by a pair of PZT crystals mounted in pockets on the large diameter portion of the resonator. The pockets are located on 90° faces of the resonator. Expanding and contracting the PZT crystals along the length of the resonator by applying bipolar voltages causes the resonator to flex in the first bending mode in two perpendicular directions. By manipulating the relative phase between the driving signals of the two PZT crystals, the




Figure 2. LLNL's diamond tool resonator.

motion of the tool can be varied from almost linear to elliptical.

The magnitude of the displacement of the tool tip is a function of the driving PZT frequencies and phase. A Michelson interferometer was used to measure the peak-to-peak motion amplitude of a tool in a single direction. Data were obtained by varying the PZT frequency and measuring the peak-to-peak displacement. **Figure 3a** shows the displacement data as a function of driving frequency and **Fig. 3b** is the same data, but from 0 to 10 kHz. The largest displacements are more than 45 nm and occur at about 2.5 kHz and 4.0 kHz. Note that these

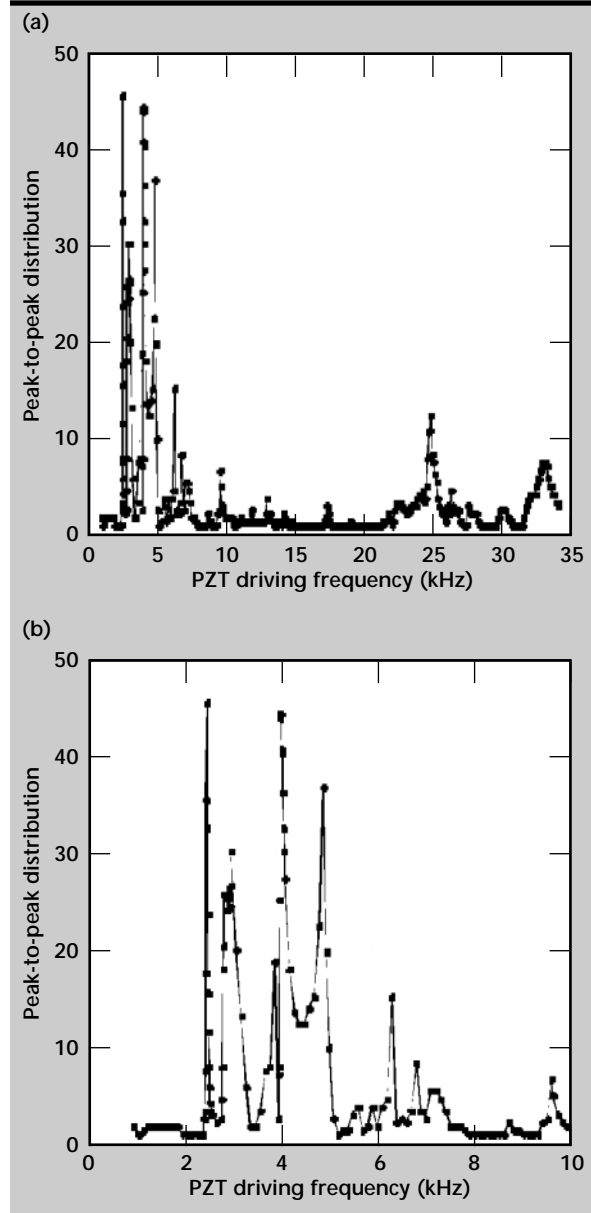


Figure 3. Single direction, peak-to-peak tool displacement data.

frequencies are below the ultrasonic frequencies. Further analysis and testing is being conducted to investigate the reason these frequencies generate the largest displacements in the resonator.

Figures 4a and 4b are Wyko white-light surface interferograms of a copper workpiece diamond-turned to examine the effect of applying elliptical tool vibration (2.54 kHz) to the diamond tool. **Figure 4a** is a diamond-turned surface without tool vibration. The striations from the tool path are clearly seen. The copper surface topography changes dramatically in **Fig. 4b** when the tool vibration is introduced. Both surfaces used identical process parameters (0.76 m/s, 5.08 μm depth of cut), except that the surface in **Fig. 4b** used the tool vibration. As expected, the surface roughness with the vibrating tool is higher than without the tool vibration.

Experimental Diamond Fly Cutting

Fly cutting is a popular machining technique that can be used to generate precision flat surfaces using single-point diamond tooling. A specialized, precision fly cutting apparatus was developed at Pennsylvania State University (PSU) under Prof. E. Marsh to investigate possible improved tool

wear using this form of intermittent contact machining with brittle material workpieces. **Figure 5** shows the precision machining hardware that consists of two air-bearing spindles mounted on a Moore No. 3 Jig-Grinder machine base.

The single-point diamond tool is mounted in a fixture on the face of one of the spindles while the other spindle is used to hold the workpiece (in this case, a silicon workpiece). The workpiece spindle is held stationary with the silicon part mounted on the face, while the tool spindle rotates the diamond tool. Fine in-feed motion control is generated by increasing the air pressure in the air bearing of the workpiece spindle, which produces 11 nm of axial growth per kPa of increased air pressure. This fine in-feed capability provides the ability to fly cut with submicron depths of cut.

Figure 6 shows an example silicon workpiece that has been fly cut with a single-point diamond tool. The overall workpiece surface roughness is about 17 nm rms, with smaller sections having roughness values of less than 2 nm rms. Note that these data are from




Figure 4. Wyko white-light surface interferograms of a copper workpiece. (a) Vibration off, roughness 6.1 nm rms; (b) vibration on, roughness 12.1 nm rms.




Figure 5. Dual spindle diamond fly cutting set-up.




Figure 6. Fly cut silicon workpiece.




Figure 7. Atomic force microscope image of a fly cut silicon workpiece.

demonstration runs and the process continues to be optimized. **Figure 7** is an atomic force microscope image of a silicon fly cut surface from this process. There are areas that appear to be machined in a ductile mode and areas that exhibit characteristics of brittle fracture mode machining with material pull outs. Also present are streaks that are left by imperfections of the diamond tool. A goal of this project is to machine a surface that is free of brittle fractures and machine tool marks, while maintaining tool wear within acceptable levels.

Future Work

The anticipated future work for this project includes continued exploration of both the vibration-assisted single-point diamond turning and diamond tool fly cutting at PSU. Among the areas of investigation for the vibration-assisted diamond turning are optimizing the PZT ceramic driving frequencies and phases, mapping tool wear and testing materials in addition to silicon. Future work for the fly cutting experimentation at PSU will include diamond tool wear mapping, surface and subsurface damage analysis and process optimization. The results of both experimental projects will be used to support the hypothesis for reduced tool wear due to lower temperatures generated at the tool and work-piece interface.

References

1. E. Marsh (1998), "Preliminary results from flycutting of BK-7 glass," work in progress.
2. Y. Takeuchi, K. Sawada, and T. Sata (1996), "Ultraprecision 3D micromachining of glass," *CIRP Ann. Manufacturing Technology*, Vol. **45**, No. 1, January.
3. E. Shamoto and T. Moriwaki (1999), "Ultraprecision diamond cutting of hardened steel by applying elliptical vibration cutting," *CIRP Ann. Manufacturing Technology*, Vol. **48**, January.
4. T. Moriwaki and E. Shamoto (1995), "Ultrasonic elliptical vibration cutting," *CIRP Ann. Manufacturing Technology*, Vol. **44**, January.



Spatial-Frequency-Domain Approach to Designing Precision Machine Tools

Debra Krulewich Born
Manufacturing and Materials Engineering Division
Mechanical Engineering

We have developed an error budgeting methodology for the design of machines used to manufacture parts with spatial-frequency-based specifications.

Introduction

The aim of this project is to develop a methodology to design machines for the manufacture of parts with spatial-frequency-based specifications.

One of our responsibilities is to design machine tools that can produce advanced optical and weapons systems. Recently, many of the component tolerances have been specified in terms of the spatial frequency content of surface profile, waviness, and finish errors.

During the design of a new machine tool, we use an error budget as a sensitivity analysis tool to ensure that the parts will meet the specified tolerances. Error budgets provide the formalism whereby we account for all sources of uncertainty in a process and sum them to arrive at a net prediction of how "precisely" a manufactured component will likely meet a target specification. Using the error budget, we are able to minimize risk during the design stage by ensuring that the machine will produce components that meet specifications before the machine is actually built.

Minimizing the risk while maintaining accuracy is a key manufacturing goal for programs that cannot tolerate yield factors less than 100%, such as in fabricating components for the nuclear weapons program. However, the current error budgeting procedure provides no formal mechanism for designing machines that will be required to produce parts with spatial-frequency-based specifications. We have received specifications for advanced optical and weapons systems that are posed in terms of the continuous spatial frequency spectrum of the surface errors of the machined part.¹ Based on these requirements, it is no longer acceptable to specify tolerances in terms of a single number that spans all temporal and spatial frequencies. During

this fiscal year, we have developed a new error budgeting methodology to aid in the design of new machines used to manufacture parts with spatial-frequency-based specifications.

Progress

Technical Approach

Figure 1 shows flowcharts for both the conventional and the new error budget procedures. The upper portion of **Fig. 1** shows Donaldson's flowchart² and illustrates the mapping of error sources onto part geometry. The first step of the conventional error budget is to identify the physical influences that generate the dimensional errors that propagate through the machine tool. These include effects such as thermal gradients and temperature variability, bearing noise, fluid turbulence in cooling passages, and geometry errors such as linear bearing non-straightness.

The next step is to determine how each source couples to the machine. A coupling mechanism converts these physical influences into tool or part displacements (errors). An example of a coupling mechanism is the thermal expansion that may transform a time-varying heat source into a machine axis distortion. These displacements represent dimensional changes in the system. A single peak-to-valley number is usually used to quantify the dimensional changes, not differentiating between the spatial frequency content of the error. The next step is to sum all the contributing errors using an appropriate combinatorial algorithm.³ The last step in the error budgeting procedure is to transform these errors into the workpiece coordinate system.

To convert these machine displacements into the errors on the workpiece surface in the directions of

interest, we must consider the tool path (for example, feed rates and spindle speeds). The output from this procedure is a single number predicting the net error on a machined workpiece. We then compare this number to the part specifications. If the prediction does not meet specifications, we evaluate methods to improve this design by observing which sources are the dominating error contributors. In this way we can evaluate the cost versus accuracy improvement of different candidate designs.

If improvements can be made to an existing design, the error budget can be used to predict the effect on part accuracy. If modifications are not practical, we then consider an entirely new design or possibly reevaluate the specifications.

The lower portion of **Fig. 1** shows the new error budget approach. The first two steps, identifying the sources and how they couple to the machine, are identical and are explained in the previous paragraph. However, the new approach differs in the next step, where the elemental errors are converted into the frequency domain.

The next step is to combine the errors in the frequency domain. The combinatorial rule is a new algorithm with a statistical foundation. The combined errors then go through a material removal transfer function that accounts for the dynamics of the material removal process and machine tool.

The material removal process creates cutting forces. The ratio between the cutting force and amount of material removed is the material removal transfer function. These cutting forces combine with forces induced by the machine kinematic errors. The machine structure responds to these forces with dynamic tool-to-part relative displacements that ultimately result in additional errors on the machined part. The last step is to transform the errors into the

part coordinate system. The output from this process is the continuous spectrum of errors at all spatial frequencies on the part. Details on the technical approach appear elsewhere.⁴

Experimental Results

To verify the procedure, we machined copper and aluminum parts on a T-base lathe and measured the spatial frequency content of the resulting surfaces. We machined disks in a facing operation. During the facing operation, the tool shank is parallel to the axis of rotation of the spindle and feeds inward toward the centerline of the part as the spindle rotates. The final part is a disk with a nominal flat surface perpendicular to the spindle axis of rotation.

We then identified six main contributing errors: (1) spindle axial errors, (2) x-axis straightness in the z-direction; (3) squareness of the x-axis to the spindle axis of rotation; (4) spindle thermal growth; (5) thermal growth due to cycling of the environmental controller; and (6) diurnal cycles due to temperature changes during the day.

We measured the machine tool-to-part displacement errors individually and converted them to the frequency domain. We then used the combinatorial rule and mapped the errors to the surface to produce a prediction of the frequency content of the resulting surface. We compared this prediction to the measured surface along a radial trace across the part.

Results are reported for the following machining conditions:

Material: copper

Spindle speed: 180 RPM

Feed rate: 0.08 in./min

Tool radius: 0.03 in.

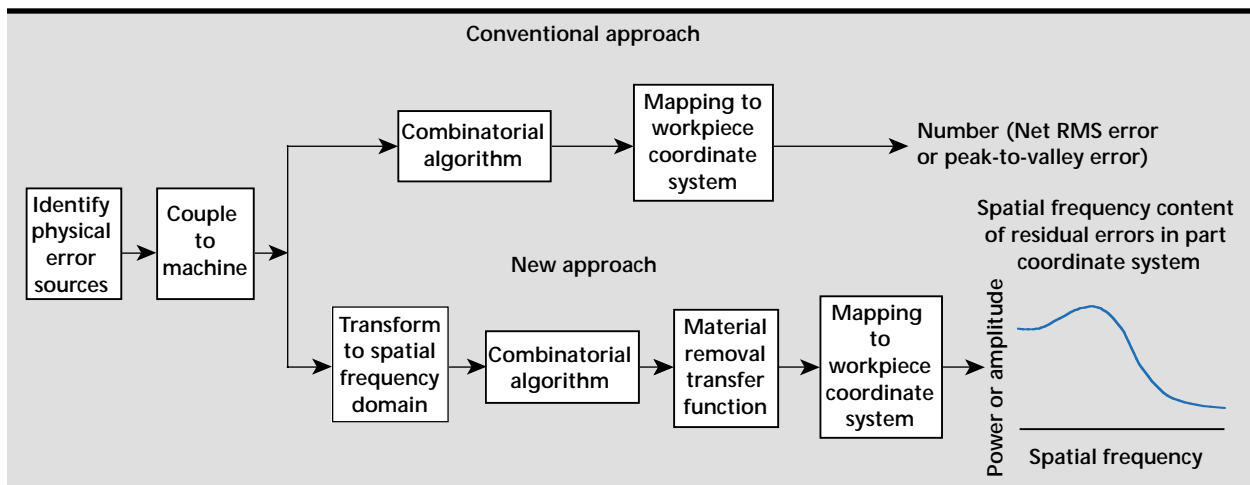


Figure 1. Flowcharts for both the conventional and the new error budget procedures.

The conventional error budget predictions of errors along a radial trace are shown in **Table 1**.

The total error across the entire part while machining copper at 180 RPM and a feed of 0.08 in./min was measured to be 13.36×10^{-6} in. This value falls within the RMS and algebraic sum as reported in **Table 1**. However, this value provides no information about the frequency content of this total error.

The results of the new error budget procedure are plotted in **Figs. 2 and 3**. **Figure 2** shows the dominant errors that occurred at relatively low frequencies. The expected errors and 95% confidence limit are arrived at through the statistically based combinatorial rule and error budget procedure. We measured errors on the machined part using a profilometer. The measured errors stayed below the upper limit with the exception of one point at approximately 1 cycle/in. This may be due to a thermal condition that existed during either the machining process or the measurement process that was not accounted for. It may also be due to some

simplifications that were made while mapping the tool geometry onto the part.

We also predicted a peak at around 2 cycles/in. that we didn't significantly observe on the measured part. This peak is due to the cycling of the environmental control. While machining, we used cutting lubricant that likely supplied thermal control at the tool/part interface, minimizing the thermal growth due to environmental cycling.

The high-frequency errors are dominated by the tool marks. During this project we did not evaluate the surface of a sharp tool or the surface of a worn tool. However, we observed that a sharp tool leaves behind not only the nominal shape of the tool, but higher frequency errors that can be approximated by a white noise of amplitude 0.01×10^{-6} in. at all frequencies. This accounts for the upper bound in **Fig. 2**. We measured errors on the machined part using an atomic force microscope (AFM). The measured errors stayed below the upper bound. Note that to predict the effect of a worn tool, the expected tool edge geometry would have to be put into the model.

Table 1. Conventional error budget results.

Parameter	Predicted error (in.)
Spindle growth	27.26×10^{-6}
Environmental control	14.00×10^{-6}
Straightness	4.50×10^{-6}
Squareness	9.70×10^{-6}
Spindle axial errors	0.51×10^{-6}
Tool marks	0.41×10^{-6}
Algebraic sum	56.38×10^{-6}
RMS	9.38×10^{-6}
Mean	32.88×10^{-6}




Figure 2. Low-frequency errors.

Future Work

Further work needs to be performed to better measure surfaces in the frequency domain. As we see more parts specified in the frequency domain, this characterization will become more important whether or not this error budgeting procedure is used. Currently we have a variety of measurement instruments, each of which has a specific resolution, range, and frequency response, and it is not clear how to combine each of these instruments to yield a continuous profile measurement over a broad frequency band.




Figure 3. High-frequency errors.


For example, the AFM can only measure the high-frequency content of a surface due to its limited range of motion, while a profilometer has the ability to measure only lower frequency components. However, the output is filtered mechanically by the shape of the stylus and cannot measure very high frequency components. A procedure that can accurately combine these measurements over a broad frequency band would be extremely useful. Without this, we have no way of verifying the accuracy of parts that have been specified in the spatial frequency domain.

For this procedure to work, we must have a method of estimating the expected errors in the frequency domain associated with each machine component. It is not practical to perform detailed tests on every machine as we did in our test study. Furthermore, in some cases, the machine may not yet exist.

Using the current error budgeting procedure, where possible analysis is performed to determine a single peak-to-valley or RMS number, at times it is possible to use the manufacturers' specifications. Where analysis is not possible and no manufacturer's data exists, we rely on prior experience and data as well as expert opinions. This task is much more complex when we must predict the function of these errors over the frequency domain rather than predicting a single number as in the conventional approach. In some cases, we may be able to perform

appropriate analysis, such as predicting the machine's dynamic characteristics. It is very unlikely that manufacturers will provide useful information since the component errors are not typically expressed in this way. Furthermore, most experts will not be familiar with characterizing errors in the frequency domain, so we cannot rely on expert advice. As more machines are characterized in this way, we hope to develop an archive of typical error types. This archive can then be referenced during the error budgeting procedure.

References

1. D. M. Aikens, C. R. Wolfe, and J. K. Lawson (1999), "The use of power spectral density (PSD) functions in specifying optics for the National Ignition Facility," *SPIE*, Vol. **2576**, pp. 281–292.
2. R. R. Donaldson (1980), "Error budgets," in *Machine Tool Accuracy*, Vol. 5 of *Technology of Machine Tools: A Survey of the State of the Art by the Machine Tool Task Force*, R. Hocken, ed., ch. 9.
3. Y. L. Shen and N. A. Duffie (1993), "Comparison of combinatorial rules for machine error budgets," *Ann. CIRP*, Vol. **42** (No. 1), pp. 619–621.
4. D. A. Krulewich (1998), "A spatial-frequency-domain approach to designing precision machine tools," *Engineering Research, Development and Technology*, Lawrence Livermore National Laboratory, Livermore, California (UCRL 53868-98), pp. 3-3–3-6. 

Power Spectral Density Measurements of Machined Surfaces over an Extended Range of Spatial Frequencies

Jon M. Baldwin
Materials and Manufacturing Engineering Division
Mechanical Engineering

We have measured the surface error of a workpiece using several instruments, then combined the data into one set. We have resolved problems concerning normalization of the data among the instruments, transfer functions of the instruments, and other power spectral density issues.

Introduction

It is becoming increasingly common to specify the quality of precision-machined surfaces, particularly those intended for optical applications, by the power spectral density (PSD) function of the dimensional error of the surface. The PSD is useful in such applications because, when captured over a suitable range of spatial frequency, it efficiently encapsulates a great deal of information about the performance of the surface.

The goal of this project is to measure the surface error of a workpiece and display its PSD over a wide spectrum of spatial frequency. This requires measuring the workpiece on several instruments and combining the data into a single set. Among the problems are: normalizing the instruments with respect to each other, assessing and incorporating the transfer function or the bandwidth of the instruments, dealing with the loss of signal to noise at high spatial frequency, and assessing the PSD confidence interval across the multiple instruments.

Use of the 1-D PSD is fairly well understood. The 2-D PSD, however, is more powerful in these instances. Integration of the 2-D PSD over a suitable range of spatial frequencies yields, directly, the RMS roughness. The 2-D PSD has been used in optical component evaluation by a number of researchers^{1,2} who have applied it to specification of components for the National Ignition Facility (NIF), and have been concerned with fabrication of optics for extreme ultraviolet lithography (EUVL). For optical components, a spatial frequency range from $1/(\text{clear aperture})$ to $50 \mu\text{m}^{-1}$ is of interest.² The metrology instruments used in this study span the approximate spatial frequency range from 0.01 mm^{-1} to $10 \mu\text{m}^{-1}$.

Some issues remain unsettled for 2-D PSDs as tools for evaluation of machined surfaces. The primary concern has been with surfaces produced by methods that are expected to generate spatially isotropic PSDs. Here, we are concerned with general surfaces including those manufactured by diamond turning, that is, surfaces with obvious lay.

The available publications provide little information about the details of data collection, data analysis, and normalization of data taken with several instruments over a wide range of spatial frequency. An example of an average radial 2-D PSD of a polished surface with reasonable continuity over six orders of magnitude of spatial frequency has been demonstrated, but with no description of steps involved, with evidence that the 2-D PSD inherently captures information regarding lay. This topic was not developed.

Progress

The purpose of our work is to reduce the PSD technique to routine practice for diamond turned surfaces, to evaluate, and, if possible, to resolve problems associated with combining data from several instruments. The specific instruments addressed here are: a Digital Instruments Nanoscope atomic force microscope (AFM), a Wyko NT2000 interferometric microscope, and a Zygo Mark GPI 4-in. interferometer. These instruments together can cover a spatial frequency range of more than six orders of magnitude, with substantial frequency overlap between instruments.

All of the above instruments are provided with software systems that provide significant data analysis power. In all cases, the details of the embedded

algorithms are unavailable. To eliminate unknown processing operations and algorithms as a source of uncertainty, only the unprocessed data from each instrument were recorded. The only operations performed with software provided by the instrument manufacturer were those necessary for conversion of data file and disk formats to be compatible with the Windows NT environment. MATLAB routines were written to read the raw data files produced by each instrument and to convert them to a common format. From this point on, all data sets were processed identically by appropriate MATLAB programs.

For each instrument/artifact/resolution combination, multiple data sets (~20) are required to (1) improve the statistical uncertainty in the PSD estimate, and (2) give a valid representation of the surface which is, at best, sampled over a small percentage of its total area by the higher resolution instruments. Each instrument must be applied at two or three resolution settings to get adequate overlap across the six-order-of-magnitude frequency range. For each measurement (except the full aperture measurements with the 4-in. interferometer) the specimen was oriented so that the x-axis was parallel to the direction of the tool feed. Each data sample was treated separately by first subtracting the least-squares plane and computing the 2-D PSD by the periodogram method, with appropriate scaling. All the PSDs for each set of conditions were averaged and a variety of display methods was used to view the individual and combined results.

A 3-D plot of the PSD works well for displaying results from a single instrument and resolution setting. A typical example, obtained with the AFM, is shown in **Fig. 1**. The x scan direction is nominally parallel to the tool feed direction, and the nature of the lay is evident. The average radial PSD used elsewhere works well for




Figure 1. Three-dimensional plot of power spectral density for single instrument.

anisotropic surfaces and is a useful intermediate in getting RMS roughness numbers, but suppresses data about lay and does not seem particularly useful in visualizing anisotropic surfaces.

For this application, two separate 2-D plots of the PSD seem useful, sectioned parallel to or across the lay direction, with the section chosen to give the maximum area under the PSD. This gives, in a sense, a worst-case estimate of the surface and preserves lay information (see **Fig. 2**). Plots thus taken show, on average, agreement between instruments and resolutions within about an order of magnitude, and are fairly linear on a log-log scale over the frequency range. An example, showing data from all three instruments at two resolution settings for each, is presented in **Fig. 2**. Agreement between instruments and instrument resolution settings needs to be improved.

A model for the instrument transfer function (ITF) of the Wyko exists³ and appears to slightly improve agreement between different resolutions of that instrument⁴ and between it and the AFM. It fails visibly and predictably, though, at frequencies greater than half the Nyquist frequency since the correction term, which appears as a divisor to the PSD, goes to zero too rapidly. Compensation for the ITF is an area that needs further and careful study. Either additional, better models should be developed, or the ITF in each case measured.

The Zygo may be particularly challenging in this respect. Zygo has both modeled and measured the ITF of a similar instrument and, while their results do not apply directly to the present situation, their methods may apply.

The other major area that needs further work is development of reasonable estimates of uncertainty of the PSD. Current estimates consider only the





Figure 2. Two-dimensional plots of power spectral density for several instruments.

statistical properties of the PSD, under somewhat unreasonable assumptions. When applied to the data of the current study, they produce estimates that, clearly, are overly optimistic.

Additional measurements are under way at lower resolution for both the interferometric microscope and the AFM. In the former case, the objective is to increase overlap with the high frequency end of the 4-in. interferometer data. In the latter case, overlap is already adequate, but it is hoped that we will be able to resolve some questions about the ITF model for the Wyko.

Finally, the user interface for the MATLAB data analysis routines needs to be improved. The current version uses a DOS-style command line interface, presently appropriate in view of the developmental status of the work. When the analysis has been reduced to routine practice, a Windows-style interface should be considered.

References

1. J. K. Lawson, C. R. Wolfe, K. R. Manes, J. B. Trenholme, D. M. Aikens, and R. E. English, Jr. (1995), "Specification of optical components using the power spectral density function," *SPIE Proceedings*, Vol. **2536**, pp. 38-50.
2. J. S. Taylor, G. E. Sommargren, D. W. Sweeney, R. M. Hudyma, and E. M. Gullikson (1998), "The fabrication and testing of optics for EUV projection lithography," Lawrence Livermore National Laboratory, Livermore, California (UCRL-JC-128290).
3. E. L. Church, T. V. Vorburger, and J. C. Wyant (1985), "Direct comparison of mechanical and optical measurements of the finish of precision machined optical surfaces," *Opt. Eng.*, Vol. **24**(3), pp. 388-395.
4. J. A. Soobitsky, L. Deck, J. Biegen, J. Kramer, K. Khalsa, T. L. Miles, and F. Demarest (1997), "Zygo engineering report: NIF high resolution large aperture interferometer," June 23. 

Dimensional Metrology of Non-Rigid Objects

Kenneth L. Blaedel and David W. Swift
Manufacturing and Materials Engineering Division
Mechanical Engineering

We have developed a method of measuring non-rigid objects for inspecting component parts. Our method of metrology also allows prediction of the shape of an assembly of non-rigid parts.

Introduction

The goal of this project is to devise and prove a method of dimensionally measuring non-rigid objects for inspecting component parts, which will further allow predicting the shape of an assembly of multiple non-rigid components. This is beyond the ability of conventional dimensional metrology, which assumes the artifact to be measured is rigid. The approach we developed is to perform a measurement while the non-rigid body is constrained under well-defined conditions. Through simulation, the effect of the constraints is then “backed out” to reveal information about the non-rigid part only. Then, the part can be constrained by a “virtual” fixture that is embodied in software to determine whether the object meets its specification.

The virtual fixture must be designed to reveal those aspects of the part that are functionally important to its end use. While using this approach may appear to be a straightforward analysis problem, it is subject to many sources of error that must be quantified and reduced to apply such analysis with high precision.

No method of dimensional metrology exists that adequately characterizes non-rigid parts to high accuracy. Inspections are often made by

constraining a non-rigid part in a rigid fixture that controls the low-order, compliant modes of the part and then the part is measured as though it were rigid. Hence, the measurements inseparably confound the effects of the part itself and the effects of the fixture.

This year’s task was to devise a method that would allow the separation of the effects of the fixture, which could then lead to a decision to accept or reject a non-rigid part.

The method to determine the acceptability of a non-rigid part that resulted requires two transformations of the measurement data, as shown schematically in **Fig. 1**.

The first step in the process is to inspect the non-rigid part, for example, on a coordinate measuring machine. The data that result reflect both errors of the part and deformations of the part that are caused by the fixture that holds the part during inspection.

The second step is to perform a structural analysis to calculate the shape of the part, as though it were free from external constraints and external forces such as gravity. The predicted shape is referred to as the “free” state of the part and is not usually physically realizable. The importance of

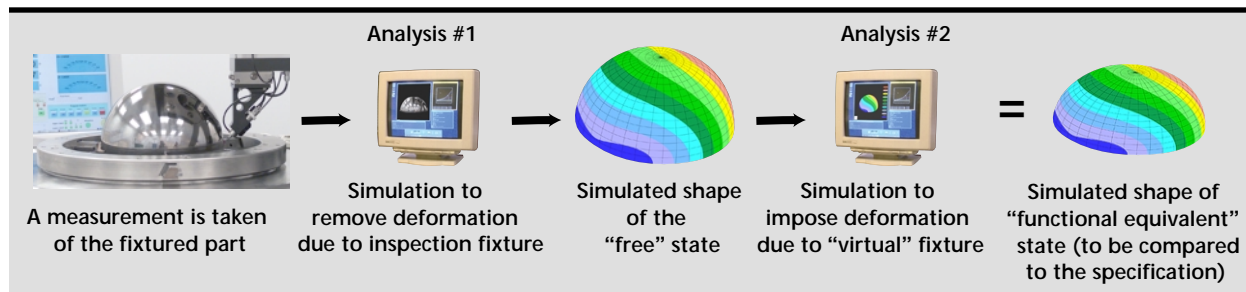


Figure 1. Methodology for dimensional inspection of a non-rigid component.

this state is that it reflects only errors in the part itself and does not include deformation induced by the fixture. This state is ordinarily not comparable to the specification of the part because for non-rigid parts, the low-order deformations, which are associated with the most compliant modes, are not the more important to the function of the part. Hence, the specification of the part often allows significant error in the low-order, compliant modes. The exercise is often to measure the higher order modes in the presence of significant low-order modes.

Therefore, before the free state can be compared to the drawing that specifies the part, the free state must have the low-order, compliant mode(s) removed. This is the purpose of the third step, which is shown in **Fig. 1** as the imposition of a “virtual” fixture. The virtual fixture is one that preserves characteristics that are important to the function of the part and removes the characteristics that are less important. Its general purpose is to simulate the functional state of the part, and results in the “functional-equivalent” state shown in **Fig. 1**.

In the fourth step, this state is directly compared to the drawing that specifies the part to determine the acceptability of the part.

In addition to enabling a method of inspection for non-rigid parts, the method of metrology proposed above also allows the prediction of the shape of an assembly of non-rigid parts. As shown in **Fig. 2**, the simulation of the assembly of two non-rigid parts, each in its free state, is performed by imposing a common boundary between them. Because current metrology methods confound deformation arising from the fixture and the part, the inspection data do not allow accurate prediction of the distortion that will occur when two non-rigid components are assembled.

Progress

Three particular tasks require further development to make this lengthy method of metrology successful. First, the uncertainty introduced in the measurement from the finite-element analysis (FEA) needs to be assessed. Second, fixturing needs to be devised that particularly allows its effect on the deformation of the part to be accurately modeled. Third, virtual fixturing needs to be conceptualized that enables the evaluation of important properties of the part and “filters out” those properties that are not. Of these, the assessment of uncertainty in performing an FEA is being pursued first.

Dimensional Uncertainty Introduced by Structural Analysis

Modern dimensional metrology requires that a statement of uncertainty accompany all dimensional measurements.¹ The basis for assessing the overall uncertainty is a comprehensive error budget which identifies each potential source of error, quantifies its probable magnitude, assigns the error to a work-piece category, and then combines the individual workpiece errors to estimate the total uncertainty associated with the measurement.

For the integration of dimensional measurement and analytical modeling, a simulated deformation that is caused by the fixturing constraints of a non-rigid part is one component of the best estimate of its real geometry (such as its dimensions). The error associated with analytical prediction is unusual to the practice of dimensional metrology and its uncertainty must be assessed to make a statement about the overall dimensional uncertainty.

The uncertainty of a simulated deformation is also established through an uncertainty budget and

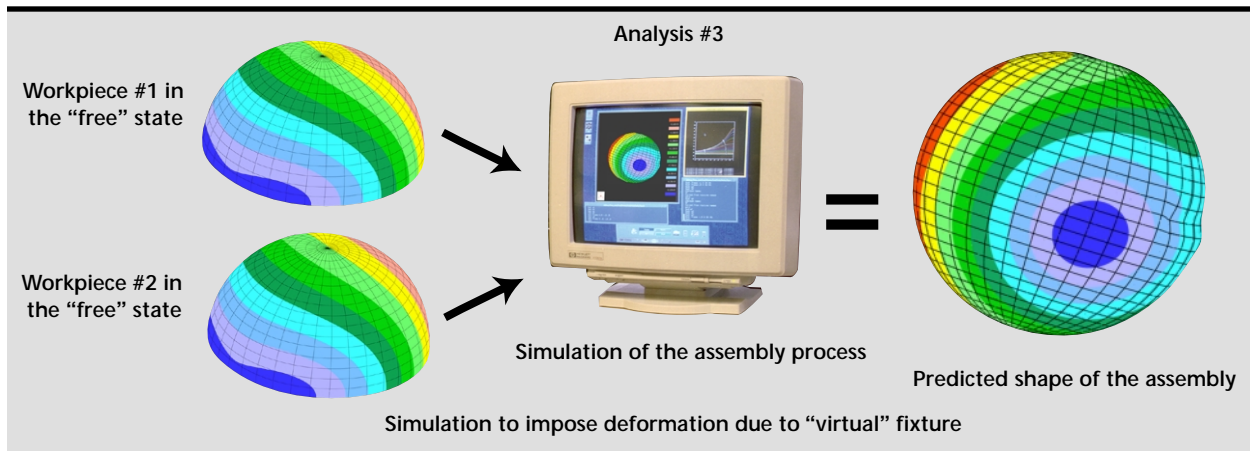


Figure 2. Methodology for simulating the assembly of two non-rigid bodies.

is corroborated by comparing the deformation calculated from simulation to that measured in a “verification” experiment. The verification geometry chosen for the initial experiment is a thin right-circular cylinder because its analysis is straightforward and easy to measure interferometrically. The initial method of simulation will be a linear finite-element method because of its generality. There are three general categories of errors to quantify.

The first category of error (*validation*) is associated with the overall accuracy of the simulation and how quickly it converges to the correct solution for the given boundary conditions. This will be assessed by conventional methods, such as comparing the difference in calculated deformation as a function of discretization error, element kinematics and material inhomogeneity. Also, for the geometry of the verification experiments, the accuracy of the simulations will be compared to an analytical solution.

The second category of error (*verification*) concerns the difference between the deformation predicted by the simulation and the deformation measured in a verification experiment. A dominant source of error in this category is the modeling of the material and determining its properties. Because material properties given in handbooks can easily be in error, the important properties, such as the elastic moduli and Poisson ratios, will have to be measured for the actual material used in the experiment.

In addition, the homogeneity of the material may have to be quantified for the actual experiment. Other material properties such as creep and plasticity will be assessed but will not likely be required in the simulation. These effects will be minimized in the validation experiment by selecting a material known to be elastic.

A second significant error in corroborating calculation with experiment is in the modeling of constraints, for example, the “contact algorithms.” It is not known in advance whether this will be a significant error contributor, so the validation experiments will yield the first estimates of the error. Some iteration on the contact algorithms may be required to determine the best contact algorithms to use for thin shells.

The initial verification experiments are devised to check exactly this category of error. They will focus on a thin cylindrical shell where deformation can be calculated by closed-form analytic solution. A cylinder is also easy to mesh and inexpensive to measure.

The third category of error (*application to a real problem*) is to estimate how well the simulation

will perform under different, real conditions beyond those of the verification experiment, and in particular, for a different geometry. This category of error is somewhat less studied than the first two categories.

Verification Experiment

The geometry for the verification experiment is a thin cylinder whose wall thickness-to-radius ratio is about 1:100 and governs how “non-rigid” the part is. Its length-to-radius ratio is 1:1 so that radial deformation at one end of the cylinder is quite different from the deformation imposed by a constraint at the opposite end of the cylinder.

This geometry was chosen because it is well suited to validation and verification. Easy validation arises because the deformation of a cylinder can be calculated by closed-form analytic solution, against which the results of the finite-element code can be directly compared. Easy verification arises because the deformation of the cylinder can be quickly and accurately measured for comparison to the analyses.

The status of the verification experiment is that the cylinder is being fabricated as a free-standing, electro-formed nickel shell. The electro-formed material is known to be quite homogeneous. The residual stresses induced by the electro-forming process, however, will cause the shell to be non-round, which is the desired effect. The measurement will be performed with a cylindrical, grazing-incidence interferometer,² which gives a very quick characterization of the deformation of the entire cylindrical surface.

Figure 3 shows the structural analysis of the verification cylinder supported by three wires, that is, three constraints that are each very compliant in five degrees of freedom and very stiff in the sixth degree of freedom. This may not be the best method of fixing the cylinder during measurement. The analysis calculates deformation for a radial load that is diametrically imposed on the end of the cylinder opposite from the three constraints. The verification experiment involves taking measurements of the unloaded cylinder, which provides the measurement data in **Fig. 1** and then imposing the radial load to compare the measured deformation directly to the analysis. The actual material properties of the electro-formed shell will be measured by conventional methods.




Figure 3. Structural analysis of verification cylinder. A radial load imposes a complex pattern of deformation on a thin cylindrical shell.

Future Work

This past year's work was the formulation of a method of metrology for non-rigid parts. Next year's tasks are to actually qualify, by the method described above, the use of a structural analysis code by quantifying the uncertainty with which it predicts the deformation.

Also for next year is to determine how to fixture non-rigid components for inspection that particularly allows its effect on the deformation of the part to be accurately modeled.

Current metrology methods inspect non-rigid components under over-constrained and often ill-defined conditions. Such inspections suffer from two shortcomings. First, they are not traceable from one institution to another unless the two institutions agree to the use of identical inspection fixtures. Second, such data do not allow accurate prediction of the distortion that will occur when two non-rigid components are assembled. The approach to this problem is to develop classes of restraints/fixturing for components during inspection that fall into the category of "exact-constraint" and that allow the accurate simulation of their effect on the deformation of the object.

Beyond next year, we will develop the concept of virtual fixturing, which allows the evaluation of important properties of the part and "filters out" those properties that are not. Specific application will be for thin cylindrical and spherical shells. This problem is largely unexplored.

References

1. "Geometrical Product Specification (GPS)—Inspection by measurement of workpieces and measuring equipment," (ISO 14253:1997).
2. Tropel "CylinderMaster 25" available from Carl Zeiss, 7008 Northland Drive, Minneapolis, MN 55428. 